



U.S. Department of Energy Smart Grid Investment Grant

Technical Advisory Group Guidance Document #5^{*}

Topic: Techniques for Estimating Impact Measurements

August 30, 2010

OBJECTIVE

This guidance document describes the analysis techniques that are appropriate for estimating the impact metrics that are to be reported in describing the results of Consumer Behavior Studies (CBS) conducted in conjunction with Smart Grid Investment Grants (SGIG). The document:

- Describes the customer usage impact metrics that are to be calculated
- Details the treatment variations that will occur in the CBS experiments
- Details types of control groups that may be developed as part of the analysis
- Summarizes analysis techniques and their strengths and weaknesses
- Provides recommendations on the preferred combinations of control groups and analysis techniques

^{*} The following individuals on the Lawrence Berkeley National Laboratory Technical Advisory Group (TAG) drafted and/or provided input and comments on one or more of the U.S. Department of Energy Smart Grid Investment Grant (SGIG) Technical Advisory Group Guidance Documents: Peter Cappers, Andrew Satchwell and Charles Goldman (LBNL), Karen Herter (Herter Energy Research Solutions, Inc.), Roger Levy (Levy Associates), Theresa Flaim (Energy Resource Economics, LLC), Rich Scheer (Scheer Ventures, LLC), Lisa Schwartz (Regulatory Assistance Project), Richard Feinberg (Purdue University), Catherine Wolfram, Lucas Davis and Meredith Fowle (University of California at Berkeley), Miriam Goldberg, Curt Puckett and Roger Wright (KEMA), Ahmad Faruqui, Sanem Sergici, and Ryan Hledik (Brattle Group), Michael Sullivan, Matt Mercurio, Michael Perry, Josh Bode, and Stephen George (Freeman, Sullivan & Company). In addition to the TAG members listed above, Bernie Neenan and Chris Holmes of the Electric Power Research Institute also provided comments.



REQUIRED IMPACT METRICS

Table D-1 of the Guidebook for ARRA Smart Grid Program Metrics and Benefits describes eight impact metrics that are to be used to describe the impacts of pricing and information and automation/control technology on customer electricity consumption.¹ These impacts fall into two categories – impacts on electricity consumption and demand; and demand elasticities.

Both impacts on electricity use and demand elasticities should be calculated for most studies. The recommended practice involves a stepwise analysis of the impacts of pricing and/or information and automation/control technology on customer usage. In the first step, customer electricity consumption should be carefully analyzed employing analysis techniques that are appropriate for estimating the impacts of the pricing change(s) on electricity consumption *per se*. That is, the intervention effects should be directly estimated. Then, if possible, in the second step, demand models should be estimated describing the *percent change* in electricity consumption that results from the combinations of price and/or technology that are embodied in the experiment. If the results from the two analyses do not agree, then one of the two analysis models is incorrectly specified and an investigation should be undertaken to unearth the reason for the discrepancy.

Direct estimation of the impacts on electricity use provides the results *devoid* of any theoretical assumptions about how customers respond to dynamic pricing. In some rare instances, it will not be possible or appropriate to apply both modeling frameworks. For example, in situations where prices do not vary, it makes no sense to estimate load impacts using a model that is driven by pricing variation.

Electricity Consumption / Load Impact Metrics

The table contains the following metrics describing electricity consumption and demand:

1. - Annual electricity consumption – % change and MWh/year;
2. - Average hourly electricity consumption – % change MWh/hour (hours for which the calculation is required depend on specific rate design);
3. - System coincident peak demand – % change and MW (at the time of system peak); and
4. - Reliability requirements -- % change and MW (impact of planning reserve margin – specific to each locality)

It is understood that most, if not all of the consumer behavior studies that will be undertaken will rely on sampling to study the impacts for a representative cross sections of customers. Correspondingly Table D-1 also calls for estimation of sampling errors associated with the above metrics.

The electricity consumption metrics are all *raw* measurements of electricity consumption or demand. That is, they are measurements of kWh usage or kW demand. Correspondingly, they can be used to

¹ U.S. Department of Energy, "Guidebook for ARRA Smart Grid Program Metrics and Benefits," Washington, DC, December 7, 2009.



describe the changes in electricity consumption and electric demand that are observed when customers are exposed to particular combinations of pricing and/or information and automation/control technologies. They estimate the electricity use impacts attributable to the pricing and/or technology treatment, and do not make assumptions about why or how customer respond to these factors.

Demand Elasticity Metrics

In addition to the electricity consumption metrics described above, Table D-1 calls for reporting demand elasticities for the treatments under study. Demand elasticities differ from raw measurements of electricity consumption or electric demand in that they are measurements of the *percent change* in electricity demand (usually expressed as a function of a unit change in the price). Table D-1 cites four kinds of demand elasticities that may be estimated. These include:

1. - Own-price elasticity – % change in demand for electricity resulting from a 1% change in the price of electricity, controlling for other factors that may influence demand;
2. - Cross-price elasticity -- % change in demand for electricity resulting from a 1% change in the price of a related good (i.e., price per kWh during another time period);
3. - Daily price elasticity -- % change in demand for electricity for a given day resulting from a 1% change in the price for the day;
4. - Substitution elasticity -- % change in the ratio of electricity consumption in two time periods resulting from a 1% difference in price for the two time periods.

TREATMENT VARIATIONS

In general, treatment variations that will occur in CBS experiments will consist of the following types of pricing² and information and automation/control technologies³:

1. - Pricing
 - a. - TOU – time of use pricing with two or more pricing periods within a given season
 - b. - CPP – critical peak pricing in which relatively high prices are experienced during a small number of hours within a given year
 - c. - CPP/TOU – critical peak pricing overlaid on a time of use rate
 - d. - RTP – prices varying usually on a day-ahead basis by hour.
 - e. - VPP – variable peak pricing, a hybrid of TOU and RTP, where the peak period price varies daily based on system and/or market conditions.

² See Technical Advisory Group Guidance Document #5 on Rate Treatments in Consumer Behavior Study Designs for more detailed information on different pricing designs.

³ See Technical Advisory Group Guidance Document #2 on Non-Rate Treatments in Consumer Behavior Study Designs for more detailed information on these and other types of information and automation/control technologies.



- f. - PTR – peak time rebate pricing In which a customer is paid for reducing electricity below some baseline during a small number of hours within a given year
- 2. - Information Technology
 - a. - IHD – in home displays
 - b. - Web Portal – AMI meter information pushed to the web on a day late basis where it can be accessed by a variety of devices
 - c. - Push to PC/Phone – AMI meter information pushed to web in real time where it can be accessed by a variety of devices
 - d. - Bill alerts – frequent messages (usually electronic) displaying up to date cost information
- 3. - Automation/Control Technology
 - a. - HAN – home area network applications
 - b. - DLC – utility direct control of designated electricity-consuming devices

The above pricing structures are sometimes combined with the information and automation/control technologies to identify *incremental* impacts that may arise when *enhanced* information about prices is provided or when the consumers' ability to control their loads is improved. All of the information and automation/control technologies except the bill alerts should be treated as single experimental factors which do not vary over the course of the experiment. They are either on or they are off.

Some of the pricing structures, on the other hand do vary over the course of the experiments. For example, CPP rates are in effect on some days and not on others; and real time prices vary from hour to hour throughout the course of the experiment. These are more like dosage variables whose values can change over time during the course of the experimental treatment.

MODELS FOR ESTIMATING ELECTRICITY CONSUMPTION IMPACTS

All of the CBS studies are collecting panel data. That is they are collecting hourly measurements for cross sections of customers as they are exposed to experimental circumstances (i.e., various combinations of treatment factors). Some of the studies will collect pre-test data on customers under study and others will not. Most of the studies will involve tests of the separate and combined impacts of different pricing structures and technologies set forth above. Most of the studies will produce hundreds of thousands of hourly electricity consumption measurements taken under different experimental conditions (for example a sample of hourly measurements for a cross section of 400 customers for one year will produce about 3.5 million load measurements).

A variety of statistical methods can be used to estimate the impacts of experimental treatments on electricity consumption and electric demand for the CBS studies. They include:

- 1) Difference-in-differences -
- 2) ANOVA and ANCOVA -



- 3) Panel regressions
 - a. Fixed effects models
 - b. Random effect models
 - c. Instrumental variables models
- 4) Individual customer interrupted time series regressions

The above approaches are not equivalent and indeed the first two should *not* be considered for estimation of the impact metrics for CBS studies. They do not take advantage of the time series aspect of the CBS study designs; and thus produce relatively low powered comparisons (of treatment effects) that may lead to the erroneous conclusion that a treatment effect did not exist when in fact the test itself simply wasn't powerful enough to detect it. We discuss these two simple but weak models for analyzing the CBS data only to provide a context for thinking about the advantages of the more robust and precise techniques that have been developed for analyzing panel data.

Difference-in-Differences

A relatively simple model for estimating electricity consumption and load impacts requires calculating the *difference in electricity consumption or electric demand (before and after treatment)* between consumers who have been randomly assigned to a given combination of pricing and/or information and automation/control technology and a control group that was not exposed to the combination. This model can be expressed simply as follows:

$$\text{Impact} = (\text{post}_{\text{treatment}} - \text{pre}_{\text{treatment}}) - (\text{post}_{\text{control}} - \text{pre}_{\text{control}})$$

This is how most text books describe analyzing the results obtained from a RCT. With such simple models, a t-test or analysis of variance (ANOVA) can be used to describe the difference in differences and the likelihood that the observed difference that was observed could have occurred by chance alone.

This is the simplest approach to estimating the impacts of experimental treatments and it is very easy to calculate, but it is not useful in studying changes in electricity consumption or electric demand. The principal reason that this approach is unlikely to be fruitful in consumer behavior studies is that electricity consumption varies *greatly* among randomly selected consumers as a result of individual characteristics of members of the population. Factors such as dwelling size, household size, occupancy patterns, appliance holdings and other features strongly influence both electricity consumption and electric demand. Variation in these attributes in the population causes a lot of statistical noise (i.e., variation) in a simple random sample of electricity consumption. To make matters worse, the impacts of treatments on electricity consumption or electric demand are often relatively small – between 2% and 20% change. Taken together these two aspects of the problem may require *very* large sample sizes to create enough statistical power to detect meaningful differences when a simple difference of differences calculation is used to describe treatment impacts.



Analysis of Covariance

It is possible to isolate some of the noise caused by variation in the cross section using a procedure known as analysis of covariance (ANCOVA). In such a design, electricity consumption or electric demand are predicted in a regression function containing measurements of the *observable* factors that cause variation in electricity consumption within the cross section along with indicator variables representing when the treatment is present and absent. Such regression functions include unique intercept and slope parameter(s) for the treatment and control group. If the slope parameters for the treatment and control groups are not significantly different, then the difference in the intercepts can be interpreted as the treatment effect. Otherwise, the impacts of the treatment must be interpreted as an interaction between the treatment and the population characteristic(s) for which the slopes are different for the two groups. While the use of the covariates to control for variation in the cross section can improve the efficiency of the estimation of the impact on electricity consumption it fails to take advantage of significant additional information that is being collected in the panel data structures collected in CBS studies – namely the effects of variables that change over time.

Panel Regression Models

In the context of panel measurements electricity consumption can be said to vary as a result of three kinds of factors. They are:

1. - Factors that are *fixed* for any given customer over time but *vary* over customers (e.g., household income, dwelling type, dwelling size, appliance stock);
2. - Factors that do not vary over customers, but vary over time (e.g., day of week, hour of day); and
3. - Factors that vary over time but interact with factors that are fixed for a given customer (e.g., weather and central air conditioning).

Because most of the variation in electricity consumption can be explained by the above types of factors, it is almost always preferable to apply impact estimation techniques that *take account* of the impacts of these factors in estimating the impacts of price and information and automation/control technology on electricity consumption. These models will produce both more accurate and more precise estimates of treatment impacts than other simpler alternatives (i.e., differences in differences or ANCOVA designs).

Panel regression models are also referred to as longitudinal regression and time series cross-sectional models. They apply to data structures where repeated measures are collected for each unit of observation over the course of time. Hourly load data by customer is by nature a statistical panel that should be analyzed using methods designed for analyzing data that has a panel structure.

Panel regression models can explain significantly more variation in electricity consumption than other techniques because they can account for the effects of time varying factors that cause it. As such they are much more powerful than alternatives that do not take account of time varying factors. As more of the variation in electricity use patterns is explained in the panel regression model, the estimates of



treatment impacts become increasingly precise. It is also true that the more variation in electricity use that is explained by factors in the regression model, the less likely it is that omitted factors are confounded with treatment impacts. A key advantage of panel models is that depending on how they are specified, they have the ability to control for the effects of omitted variables that may otherwise be confounded with treatment impacts. In other words, they can, under some conditions, be used to estimate the impacts of treatments when control group members are different from treatment group members on *unobserved* omitted variables.

While panel regressions can increase the accuracy of impact estimates for the average customer, they cannot be used to meaningfully describe the impacts of factors that vary within the cross section. But this is not a major drawback since treatment impacts for specific groups of customers can be estimated by interacting customer or day characteristics with the variables representing the treatment impacts; or by specifying panel regression models for interesting types of customers.

The balance of this discussion focuses on three types of panel regressions; models with fixed effects (and/or time effects), models with random effects, and models with instrumental variables. There are additional variants of panel regressions that not included for brevity.⁴

Random Effects Models

In instances where a study employs well-executed randomized assignment to treatment and control groups, all of the panel regression techniques described herein will produce unbiased impact estimates. However, they are not all equally powerful. The most powerful estimation technique – the random effects model – assumes there are no omitted variables in the regression function causing the disturbance (error) term in the estimation equation to become correlated with the value of the treatment variable. Random effects models are named for the assumption that is made in estimating them – namely that all of the *omitted* variables in the regression function have random effects on the dependent variable. Correspondingly, they are assumed not to be correlated with any of the other independent variables included in the regression function. This is a very *strong* assumption *that is only justified when observations have been randomly assigned to treatment and control conditions* and no significant attrition in the treatment or control groups has occurred.

Random effect models use more information from the data structure and produce narrower standard errors. They will produce more accurate and precise estimates of treatment effects when the random effects assumption is true. If this assumption is false, these models can result in significant impact estimation errors. This model is not recommended if randomized assignment is not employed or if significant attrition from treatment or control groups has occurred. In practice, it very difficult to

⁴ A few notable mentions include difference-in-difference models with panel regressions, and Arellano-Bond estimators. Both of these methods are computationally intense given the large volume of interval data expected from the studies.



achieve the conditions required to employ the random effects model. For this reason, it is very seldom used in practice.

Fixed Effects Models

Fixed effects models relax the random effects model assumption slightly. They assume that the effects of omitted variables on the dependent variable are *fixed* within members of the cross section and within time periods. In this model, impacts of treatments are measured after removing differences between members of the cross sections and differences within members over time from the measurements. This is achieved by transforming the measurement of the dependent variable so that they are expressed as *deviations* from their within subject/time period means. In this way differences between subjects and time periods within subjects are removed from the analysis – thus *fixing* the effects of omitted variables on both dimensions.

The fixed effects models are less powerful than the random effects models but they will produce more robust estimates of program impacts when random assignment has failed in some manner. This condition is not unusual. In fact, it is more likely to have occurred than not. It can occur, for example if significant attrition from the treatment groups occurs during the study or if a significant number of customers who were assigned to an experimental treatment refuse to accept it.

The ability to control for the effects of certain kinds of omitted variables is an important strength of the fixed effects model, but it is very important to understand its limitations. It is not a magical elixir curing all kinds of selection bias. When randomized assignment is not possible or fails, it will provide more robust estimates of treatment effects than a random effects model, but it is still possible for treatment impacts to become confounded with certain kinds of omitted variables; and careful testing is required to ensure that this has not occurred. Fixed effects models are not a substitute for properly constructed control groups. They do *not* control for factors that vary over time and participants (such as occupancy patterns); or for fixed household characteristics such as air conditioning and space heating that interact with both time-varying occupancy and weather patterns. There are certain well known time varying factors that can interact with omitted variables in the cross sections of electricity consumption measurements. For example, AC ownership (which may be omitted from the cross section factors included in the regression) interacts with temperature to produce changes in hourly electricity consumption. If the prevalence of AC ownership somehow is different between the treatment and control groups, inaccurate estimates of treatment impacts can be obtained from fixed effects models even when customers are subject to the same weather patterns. Differences between the AC ownership in treatment and the control groups are not accounted for via a fixed effects model because the AC constantly interacts with weather and occupancy (which vary with time). This can lead to biased estimates of program impacts. The bottom line is that the fixed effects model is not a silver bullet for controlling for selection bias when estimating the impacts of experimental treatments on electricity consumption. When it is used, careful efforts should be made to ensure that the control and treatment groups are not different with respect to variables such as occupancy and the prevalence of air



conditioning. Evidence of differences for both factors can be observed by examining daily load shapes for treatment and control groups.

Models with Instrumental Variables

When it is not possible to achieve a completely randomized design, it is often possible to apply an experimental design referred to as a randomized encouragement design (RED). In the RED design, subjects are randomly assigned to varying levels of *encouragement* to accept the treatment. In most practical applications they are either encouraged or not encouraged. In this design, the encouragement variable serves as an instrumental variable that can be included in a regression equation to control for the effects of selection bias – the most serious threat to internal validity arising in experiments in which a RCT cannot be achieved. A valid instrumental variable only influences the dependent variable, electricity consumption, through its relationship to the treatment effect. Random assignment to the encouragement condition is critical to achieving this condition. When this occurs, the instrumental variable (encouragement) meets two critical requirements. It is correlated with the treatment impact; and it is uncorrelated with the remaining unexplained error after all other factors have been included in the prediction model. Including this variable in the regression function thus eliminates the bias that would otherwise occur in the estimation of the treatment effect.

The effectiveness of the instrument in controlling for selection bias depends on its correlation with acceptance of the treatment. An instrument is considered weak if it is not strongly correlated with acceptance of the treatment. It is possible to achieve varying levels of encouragement depending on the design of the enrollment program used in the study. An Opt-Out enrollment strategy might be thought of as a relatively *strong* instrumental variable, while an Opt-In enrollment strategy is a weak one. Because an instrument is randomly assigned, its use in the regression function estimating program impacts will result in unbiased estimates. However, if the instrument is weak it will produce relatively large standard errors and uncertainty about the magnitude of program impacts.

Individual Customer Regressions

Individual customer regressions, as their name implies, use a within-subjects design in which common regression models are used to estimate individual time series regressions. For obvious reasons, this approach is really only appropriate when a relatively large number of time periods are observed; that is the case in virtually all CBS experiments. This approach also relies on the existence of pre-treatment data and or a repeated treatment data structure. Control groups are not required for some of these designs, but individual customer regressions can be estimated for control group members if they are available.

However, if randomized assignment to treatment and control groups has been applied, more conventional techniques (panel regressions) should be employed instead. Likewise, if a control group is selected via other techniques (i.e., randomized assignment was not feasible) and is indeed comparable to the treatment group, techniques that make use of the control group should be employed. Use of



individual customer regressions is really only appropriate when no control group is available; or if there the evidence indicates the control group is not comparable.

To estimate individual customer regressions it is necessary to observe electricity consumption before and after the treatment has been applied. The results are more robust if the rate or technology treatment provides alternating or repeated treatment applications. For example critical peak pricing overlaid on a flat rate (pure CPP) and direct load control technologies typically produce repeated treatments over some experimental observation period. The treatments are introduced under varying conditions (e.g., on some days and not in others), making it possible to observe behavior with and without the treatment under similar conditions. When treatment variables are repeatedly administered it is possible to measure the extent to which the outcome—electricity consumption—rises or falls with the presence or absence of the treatment. This approach can be very effective but it only works if the effect of the event dissipates after it is removed.

When valid control groups are unavailable, individual customer regressions have several advantages, particularly in repeated treatment framework. An advantage of these models is that they can be used to describe the distribution of customer load reductions as well as the distribution of percent load reductions that occurs within a given population. That is, they can be used to show the fraction of the customers that are changing their electricity consumption to varying degrees. Often times, the information about how treatment impacts vary is just as important to policy decisions as is the accurate information about average treatment effects. By employing individual customer regressions, it is possible to control for omitted time-invariant customer characteristics and, in addition, it is possible to control for key variables such as the presence of air conditioning that are not directly observed. For example, the effect of air conditioning is captured through temperature variables and their interaction with hourly binary variables. By allowing individual customer coefficients to vary, the results are more accurate at the customer level.

MODELS FOR ESTIMATING DEMAND FUNCTIONS

Utility executives and policy makers will often want to know what would happen in the future if other prices materialize. In order to obtain such estimates, price elasticities and demand curves have to be estimated. These can be estimated in one of two ways -- first, using single equations which explain the behavior of usage as a function of price in that period and perhaps in related periods and, second, using systems of equations derived from the theory of utility maximization. In both cases, the regression models exploit the price variation over time and across customers to parameterize the demand functions. The more price variation in the models, the more precise will be the parameter estimates. Thus, the experimental design will ideally feature more than one price per pricing period. However, it may be possible to obtain price elasticities even within a single price per period since the control group will provide yet another price as will the treatment group in the pre-treatment period.



Demand equations can be estimated using the data on both treatment and control groups to predict the price elasticity of customers. If pre-treatment data are available, they can be used to adjust for any pre-existing differences between the treatment and control groups. Furthermore, weather data in conjunction with data on the presence of information or automation/control technologies (e.g., the Energy Orb, programmable communicating thermostats) and customer socio-demographic variables may also be used to explain variations in individual customers' demand for electricity.

Demand models allow for estimation of the impact of prices other than those used in the program and this is *the main strength of these models* as opposed to alternative methods such as ANOVA/ANCOVA. The transfer of the available information from existing prices to other potential prices is made possible by the use of price elasticities.

There are several types of demand elasticities. The *own price elasticity* of demand measures the percent change in demand of a good due to a one percent change in the price of the given good after controlling for all other factors that could potentially affect the demand for the good. The *cross price elasticity* measures the percent change in demand of a good due to a one percent change in the price of a related good. In single equation models, the own price and cross price elasticity are estimated separately for peak and off-peak periods.

For demand systems, a slightly different approach is used. There is one equation that measures changes in daily energy consumption and another equation that measures changes in the load shape. The *daily price elasticity* is used to measure changes in daily usage. The *elasticity of substitution*, measures the percent change in the ratio of consumption between two periods due to the change in the ratio of prices between these two periods, is used to measure changes in load shape. The two equations are jointly estimated. Predictions about demand response are made by solving the equations for values of peak and off-peak consumption.

In order to infer demand curves and elasticities, it is necessary to have price variation in the sample. The greater the variation in prices, the higher the precision in the estimated elasticities and demand curves. Ideally, one would test multiple price points for each pricing period. However, it is still possible to estimate demand price elasticities with a single pricing treatment as long as there is another price for the control group and/or the pre-treatment period. However, with just two price points, only linear demand curves can be estimated which will produce arc elasticity estimates.

Another choice that needs to be made is the procedure for econometric estimation of the elasticities after the demand model is specified. In a framework which includes a cross section of customers over time, one of the panel (or cross-sectional time-series) estimation routines can be used.⁵ Fixed effects and random effect models are two widely used panel regression estimation routines. *Fixed effects*

⁵ For more information on the panel data estimation, see Jeffrey M. Wooldridge, *Econometric Analysis of Cross Section and Panel Data*, Cambridge: Massachusetts 2002.



estimation uses a data transformation method that removes any unobserved time-invariant effect that has a potential impact on the dependent variable. This model is suitable when these unobserved time invariant effects are expected to be correlated with the other explanatory variables in the model. The alternative is the *random effects* routine which is based on the postulation that the unobserved time invariant effects are random and are not correlated with other explanatory variables of the model. Statistical tests are available for guiding the choice of the two estimation methods.

In addition to the panel estimation, the question of estimating customer-by-customer demand functions often comes up. When the model is estimated at the customer-by-customer level, the estimation sample does not constitute a panel but reduces to simple time-series estimation. The customer-specific elasticity estimation is only feasible to do if there is sufficient price variation over time on a customer-by-customer basis. For instance, if one is working with real time pricing data which features hourly price variation, the estimation of customer-specific elasticities is feasible. However, if there is only one price per period for each customer, then estimating price elasticities that are customer specific is problematic and may well be empirically impossible.

Another question deals with the choice of functional form of the demand equations. One specification which is well grounded in economic theory and which has been widely estimated in the econometric literature on time-varying prices is the constant elasticity-of-substitution (CES) model.⁶ The CES modeling system consists of several equations, all but one of which measure substitution between adjacent periods and/or hours within a day and one of which measures changes in daily energy consumption. The substitution equations capture pure changes in load shape within a day whereas the daily equation captures overall energy conservation or load building. The CES system captures the non-linearity in the relationship between demand response and dynamic prices.

Besides the CES model discussed above, other more complex options used in the literature include the Cobb-Douglas, Trans-log, Generalized Leontief (Diewert), and Generalized McFadden functional forms.⁷ The nature of the problem at hand and the policy making context will usually determine which of these widely-used functional forms are best suited for the specific application. The following criteria can be used to guide the choice of functional form:

- Parsimony in parameters: a functional form should not have numerous parameters as this will increase the likelihood of the multicollinearity problem. Moreover when the sample size is small, excess parameters imply lost degrees of freedom.
- Ease of interpretation: excessively complex functional forms may contain irregularities which may not be easily detected in the richness of parameters. Also, complex transformations may

⁶ The CES model has a strong pedigree and two of its developers went on to win the Nobel Prize in economics.

⁷ See "Production Economics: A Dual Approach to Theory and Applications," Volume 1, edited by Melvyn Fuss and Daniel McFadden, Netherlands: North-Holland Publishing Company-1978.



make it computationally difficult to derive certain parameters of interest such as elasticities of substitution.

- Computational ease: models linear in parameters have a computational cost advantage as well as a more developed statistical theory. The trade-off between the computational requirements versus statistical soundness must be carefully made.
- Interpolative robustness: within the range of the observed sample, chosen functional form should produce well-behaved and economically sound parameter estimates such as positive marginal products and negative own price elasticities.
- Extrapolative robustness: functional form should lead to sound estimates consistent with the maintained hypothesis outside the range of observed data. This criterion is particularly important for forecasting exercises.

Demand model estimation yields a comprehensive set of impact metrics. However, *this comes at a cost of requiring expertise in regressions analysis and econometrics*. Project teams should weigh the costs and benefits of each approach accordingly and select the one that conforms best to the project content and constraints.